

Optimally-Weighted Herding is Bayesian Quadrature

Ferenc Huszár

Department of Engineering
Cambridge University
fh277@cam.ac.uk

David Duvenaud

Department of Engineering
Cambridge University
dkd23@cam.ac.uk

Abstract

Herding and kernel herding are deterministic methods of choosing samples which summarise a probability distribution. A related task is choosing samples for estimating integrals using Bayesian quadrature. We show that the criterion minimised when selecting samples in kernel herding is equivalent to the posterior variance in Bayesian quadrature. We then show that sequential Bayesian quadrature can be viewed as a weighted version of kernel herding which achieves performance superior to any other weighted herding method. We demonstrate empirically a rate of convergence faster than $\mathcal{O}(1/N)$. Our results also imply an upper bound on the empirical error of the Bayesian quadrature estimate.

1 INTRODUCTION

The problem: Integrals A common problem in statistical machine learning is to compute expectations of functions over probability distributions of the form:

$$Z_{f,p} = \int f(x)p(x)dx \quad (1)$$

Examples include computing marginal distributions, making predictions marginalizing over parameters, or computing the Bayes risk in a decision problem. In this paper we assume that the distribution $p(x)$ is known in analytic form, and $f(x)$ can be evaluated at arbitrary locations.

Monte Carlo methods produce random samples from the distribution p and then approximate the integral by taking the empirical mean $\hat{Z} = \frac{1}{N} \sum_{n=1}^N f_{x_n}$ of the function evaluated at those points. This non-deterministic estimate converges at a rate $\mathcal{O}(\frac{1}{\sqrt{N}})$.




Figure 1: The first 8 samples from sequential Bayesian quadrature, versus the first 20 samples from herding. Only 8 weighted SBQ samples are needed to give an estimator with the same maximum mean discrepancy as using 20 herding samples with uniform weights. Relative sizes of samples indicate their relative weights.

When exact sampling from p is impossible or impractical, Markov chain Monte Carlo (MCMC) methods are often used. MCMC methods can be applied to almost any problem but convergence of the estimate depends on several factors and is hard to estimate (Cowles & Carlin, 1996). The focus of this paper is on quasi-Monte Carlo methods that – instead of sampling randomly – produce a set of pseudo-samples in a deterministic fashion. These methods operate by directly minimising some sort of discrepancy between the empirical distribution of pseudo-samples and the target distribution. Whenever these methods are applicable, they achieve convergence rates superior to the $\mathcal{O}(\frac{1}{\sqrt{N}})$ rate typical of random sampling.

In this paper we highlight and explore the connections between two deterministic sampling and integration methods: Bayesian quadrature (BQ) (O’Hagan, 1991; Rasmussen & Ghahramani, 2003) (also known as Bayesian Monte Carlo) and kernel herding (Chen et al., 2010). Bayesian quadrature estimates integral (1) by inferring a posterior distribution over f conditioned on the observed evaluations f_{x_n} , and then computing the posterior expectation of $Z_{f,p}$. The points where the function should be evaluated can be found via Bayesian experimental design, providing a deterministic procedure for selecting sample locations.

Herding, proposed recently by Chen et al. (2010), produces pseudosamples by minimising the discrepancy of moments between the sample set and the target distribution. Similarly to traditional Monte Carlo, an estimate is formed by taking the empirical mean over samples $\hat{Z} = \frac{1}{N} \sum_{n=1}^N f_{x_n}$. Under certain assumptions, herding has provably fast, $\mathcal{O}(\frac{1}{N})$ convergence rates in the parametric case, and has demonstrated strong empirical performance in a variety of tasks.

Summary of contributions In this paper, we make two main contributions. First, we show that the Maximum Mean Discrepancy (MMD) criterion used to choose samples in kernel herding is identical to the expected error in the estimate of the integral $Z_{f,p}$ under a Gaussian process prior for f . This expected error is the criterion being minimized when choosing samples for Bayesian quadrature. Because Bayesian quadrature assigns different weights to each of the observed function values $f(x)$, we can view Bayesian quadrature as a weighted version of kernel herding. We show that these weights are optimal in a minimax sense over all functions in the Hilbert space defined by our kernel. This implies that Bayesian quadrature dominates uniformly-weighted kernel herding and other non-optimally weighted herding in rate of convergence.

Second, we show that minimising the MMD, when using BQ weights is closely related to the sparse dictionary selection problem studied in (Krause & Cevher, 2010), and therefore is approximately submodular with respect to the samples chosen. This allows us to reason about the performance of greedy forward selection algorithms for Bayesian Quadrature. We call this greedy method Sequential Bayesian Quadrature (SBQ).

We then demonstrate empirically the relative performance of herding, i.i.d random sampling, and SBQ, and demonstrate that SBQ attains a rate of convergence faster than $\mathcal{O}(1/N)$.

2 HERDING

Herding was introduced by Welling (2009) as a method for generating pseudo-samples from a distribution in such a way that certain nonlinear moments of the sample set closely match those of the target distribution. The empirical mean $\frac{1}{N} \sum_{n=1}^N f_{x_n}$ over these pseudosamples is then used to estimate integral (1).

2.1 Maximum Mean Discrepancy

For selecting pseudosamples, herding relies on an objective based on the maximum mean discrepancy (MMD; Sriperumbudur et al., 2010). MMD measures the divergence between two distributions, p and q with respect to a class of integrand functions \mathcal{F} as follows:

$$\text{MMD}_{\mathcal{F}}(p, q) = \sup_{f \in \mathcal{F}} \left| \int f_x p(x) dx - \int f_x q(x) dx \right| \quad (2)$$

Intuitively, if two distributions are close in the MMD sense, then no matter which function f we choose from \mathcal{F} , the difference in its integral over p or q should be small. A particularly interesting case is when the function class \mathcal{F} is functions of unit norm from a reproducing kernel Hilbert space (RKHS) \mathcal{H} . In this case, the MMD between two distributions can be conveniently expressed using expectations of the associated kernel $k(x, x')$ only (Sriperumbudur et al., 2010):

$$\text{MMD}_{\mathcal{H}}^2(p, q) = \sup_{\substack{f \in \mathcal{H} \\ \|f\|_{\mathcal{H}}=1}} \left| \int f_x p(x) dx - \int f_x q(x) dx \right|^2 \quad (3)$$

$$= \|\mu_p - \mu_q\|_{\mathcal{H}}^2 \quad (4)$$

$$= \iint k(x, y) p(x) p(y) dx dy - 2 \iint k(x, y) p(x) q(y) dx dy + \iint k(x, y) q(x) q(y) dx dy, \quad (5)$$

where in the above formula $\mu_p = \int \phi(x) p(x) dx \in \mathcal{H}$ denotes the *mean element* associated with the distribution p . For characteristic kernels, such as the Gaussian kernel, the mapping between a distribution and its mean element is bijective. As a consequence $\text{MMD}_{\mathcal{H}}(p, q) = 0$ if and only if $p = q$, making it a powerful measure of divergence.

Herding uses maximum mean discrepancy to evaluate of how well the sample set $\{x_1, \dots, x_N\}$ represents the target distribution p :

$$\begin{aligned}
\epsilon_{\text{herding}}(\{x_1, \dots, x_N\}) &= \text{MMD}_{\mathcal{H}}\left(p, \frac{1}{N} \sum_{n=1}^N \delta_{x_n}\right) \\
&= \iint k(x, y) p(x) p(y) dx dy \\
&= -2 \frac{1}{N} \sum_{n=1}^N \int k(x, x_n) p(x) dx + \frac{1}{N^2} \sum_{n,m=1}^N k(x_n, x_m)
\end{aligned}
\tag{6}$$

$$\tag{7}$$

The herding procedure greedily minimizes its objective $\epsilon_{\text{herding}}(\{x_1, \dots, x_N\})$, adding pseudosamples x_n one at a time. When selecting the $n+1$ -st pseudosample:

$$\begin{aligned}
x_{n+1} &\leftarrow \underset{x \in \mathcal{X}}{\text{argmin}} \epsilon_{\text{herding}}(\{x_1, \dots, x_n, x\}) \\
&= \underset{x \in \mathcal{X}}{\text{argmax}} 2\mathbb{E}_{x' \sim p}[k(x, x')] - \frac{1}{n+1} \sum_{m=1}^n k(x, x_m),
\end{aligned}
\tag{8}$$

assuming $k(x, x) = \text{const.}$ The formula (8) admits an intuitive interpretation: the first term encourages sampling in areas with high mass under the target distribution $p(x)$. The second term discourages sampling at points close to existing samples.

Evaluating (8) requires us to compute $\mathbb{E}_{x' \sim p}[k(x, x')]$, that is to integrate the kernel against the target distribution. Throughout the paper we will assume that these integrals can be computed in closed form. Whilst the integration can indeed be carried out analytically in several cases (Song et al., 2008; Chen et al., 2010), this requirement is the most pertinent limitation on applications of kernel herding, Bayesian quadrature and related algorithms.

2.2 Complexity and Convergence Rates

Criterion (8) can be evaluated in only $\mathcal{O}(n)$ time. Adding these up for all subsequent samples, and assuming that optimisation in each step has $\mathcal{O}(1)$ complexity, producing N pseudosamples via kernel herding costs $\mathcal{O}(N^2)$ operations in total.

In finite dimensional Hilbert spaces, the herding algorithm has been shown to reduce MMD at a rate $\mathcal{O}(\frac{1}{N})$, which compares favourably with the $\mathcal{O}(\frac{1}{\sqrt{N}})$ rate obtained by non-deterministic Monte Carlo samplers. However, as pointed out by Bach et al. (2012), this fast convergence is not guaranteed in infinite dimensional Hilbert spaces, such as the RKHS corresponding to the Gaussian kernel.




Figure 2: An illustration of Bayesian Quadrature. The function $f(x)$ is sampled at a set of input locations. This induces a Gaussian process posterior distribution on f , which is integrated in closed form against the target density, $p(x)$. Since the amount of volume under f is uncertain, this gives rise to a (Gaussian) posterior distribution over $Z_{f,p}$.

3 BAYESIAN QUADRATURE

So far, we have only considered integration methods in which the integral (1) is approximated by the empirical mean of the function evaluated at some set of samples, or pseudo-samples. Equivalently, we can say that Monte Carlo and herding both assign an equal $\frac{1}{N}$ weight to each of the samples.

In (Rasmussen & Ghahramani, 2003), an alternate method is proposed: Bayesian Monte Carlo, or Bayesian quadrature (BQ). BQ puts a prior distribution on f , then estimates integral (1) by inferring a posterior distribution over the function f , conditioned on the observations $f(x_n)$ at some query points x_n . The posterior distribution over f then implies a distribution over $Z_{f,p}$. This method allows us to choose sample locations x_n in any desired manner. See Figure 2 for an illustration of Bayesian Quadrature.

3.1 BQ Estimator

Here we derive the BQ estimate of (1), after conditioning on function evaluations $f(x_1) \dots f(x_N)$, denoted as $f(X)$. The Bayesian solution implies a distribution over $Z_{f,p}$. The mean of this distribution, $\mathbb{E}[Z]$ is the optimal Bayesian estimator for a squared loss.

For simplicity, f is assigned a Gaussian process prior with kernel function k and mean 0. This assumption is very similar to the one made by kernel herding in Eqn. (7).

After conditioning on f_x , we obtain a closed-form posterior over f :

$$p(f(x_\star)|f(X)) = \mathcal{N}(f(x_\star) | \bar{f}(x_\star), \text{cov}(x_\star, x'_\star)) \quad (9)$$

where

$$\bar{f}(x_\star) = k(x_\star, X)K^{-1}f(X) \quad (10)$$

$$\text{cov}(x_\star, x'_\star) = k(x_\star, x_\star) - k(x_\star, X)K^{-1}k(X, x_\star) \quad (11)$$

and $K = k(X, X)$. Conveniently, the GP posterior allows us to compute the expectation of (1) in closed form:

$$\mathbb{E}_{\text{GP}}[Z] = \mathbb{E}_{\text{GP}}\left[\int f(x)p(x)dx\right] \quad (12)$$

$$= \iint f(x)p(f(x)|f(X))p(x)dxdf \quad (13)$$

$$= \int \bar{f}(x)p(x)dx \quad (14)$$

$$= \left[\int k(x, X)p(x)dx\right] K^{-1}f(X) \quad (15)$$

$$= \mathbf{z}^T K^{-1}f(X) \quad (16)$$

where

$$z_n = \int k(x, x_n)p(x)dx = \mathbb{E}_{x' \sim p}[k(x_n, x')]. \quad (17)$$

Conveniently, as in kernel herding, the desired expectation of $Z_{f,p}$ is simply a linear combination of observed function values $f(x)$:

$$\mathbb{E}_{\text{GP}}[Z] = \mathbf{z}^T K^{-1}f(X) \quad (18)$$

$$= \sum_n w_{\text{BQ}}^{(n)} f_{x_n} \quad (19)$$

where

$$w_{\text{BQ}}^{(n)} = \sum_m \mathbf{z}_j^T K_{nm}^{-1} \quad (20)$$

Thus, we can view the BQ estimate as a weighted version of the herding estimate. Interestingly, the weights w_{BQ} do not need to sum to 1, and are not even necessarily positive.

3.1.1 Non-normalized and Negative Weights

When weighting samples, it is often assumed, or enforced (as in [Bach et al., 2012](#); [Song et al., 2008](#)), that the weights w form a probability distribution. However, there is no technical reason for this requirement, and in fact, the optimal weights do not have this property. Figure 3 shows a representative set of 100 BQ weights chosen on samples representing the distribution in figure 1. There are several negative weights, and the sum of all weights is 0.93.

Figure 4 demonstrates that, in general, the sum of the Bayesian weights exhibits shrinkage when the number of samples is small.




Figure 3: A set of optimal weights given by BQ, after 100 SBQ samples were selected on the distribution shown in Figure 1. Note that the optimal weights are spread away from the uniform weight ($\frac{1}{N}$), and that some weights are even negative. The sum of these weights is 0.93.




Figure 4: An example of Bayesian shrinkage in the sample weights. In this example, the kernel width is approximately $1/20$ the width of the distribution being considered. Because the prior over functions is zero mean, in the small sample case the weights are shrunk towards zero. The weights given by simple Monte Carlo and herding do not exhibit shrinkage.

3.2 Optimal sampling for BQ

Bayesian quadrature provides not only a mean estimate of $Z_{f,p}$, but a full Gaussian posterior distribution. The variance of this distribution $\mathbb{V}[Z_{f,p}|f_{x_1}, \dots, f_{x_N}]$ quantifies our uncertainty in the estimate. When selecting locations to evaluate the function f , minimising the posterior variance is a sensible strategy. Below, we give a closed form formula for the posterior variance of $Z_{f,p}$, conditioned on the observations $f_{x_1} \dots f_{x_N}$, which we will denote by ϵ_{BQ}^2 . For a longer derivation, see [Rasmussen & Ghahramani \(2003\)](#).

$$\epsilon_{\text{BQ}}^2(x_1, \dots, x_N) = \mathbb{V}[Z_{f,p}|f_{x_1}, \dots, f_{x_N}] \quad (21)$$

$$= \mathbb{E}_{x, x' \sim p}[k(x, x')] - \mathbf{z}^T K^{-1} \mathbf{z}, \quad (22)$$

where $\mathbf{z}_n = \mathbb{E}_{x' \sim p}[k(x_n, x')]$ as before. Perhaps surprisingly, the posterior variance of $Z_{f,p}$ does not depend on the observed function values, only on the location x_n of samples. A similar independence is observed in other optimal experimental design problems involving Gaussian processes ([Krause et al., 2006](#)). This allows the optimal samples to be computed ahead of time, before observing any values of f at all ([Minka, 2000](#)).

We can contrast the BQ objective ϵ_{BQ}^2 in (22) to the objective being minimized in herding, $\epsilon_{\text{herding}}^2$ of equation (7). Just like $\epsilon_{\text{herding}}^2$, ϵ_{BQ}^2 expresses a trade-off between accuracy and diversity of samples. On the one hand, as samples get close to high density regions under p , the values in \mathbf{z} increase, which results in decreasing variance. On the other hand, as samples get closer to each other, eigenvalues of K increase, resulting in an increase in variance.

In a similar fashion to herding, we may use a greedy method to minimise ϵ_{BQ}^2 , adding one sample at a time. We will call this algorithm *Sequential Bayesian Quadrature* (SBQ):

$$x_{n+1} \leftarrow \underset{x \in \mathcal{X}}{\operatorname{argmin}} \epsilon_{\text{BQ}}(\{x_1, \dots, x_n, x\}) \quad (23)$$

Using incremental updates to the Cholesky factor, the criterion can be evaluated in $\mathcal{O}(n^2)$ time. Iteratively selecting N samples thus takes $\mathcal{O}(N^3)$ time, assuming optimisation can be done on $\mathcal{O}(1)$ time.

4 RELATING $\mathbb{V}[Z_{f,p}]$ TO MMD

The similarity in the behaviour of $\epsilon_{\text{herding}}^2$ and ϵ_{BQ}^2 is not a coincidence, the two quantities are closely related to each other, and to MMD.

Proposition 1. *The expected variance in the Bayesian quadrature ϵ_{BQ}^2 is the maximum mean dis-*

crepancy between the target distribution p and $q_{\text{BQ}}(x) = \sum_{n=1}^N w_{\text{BQ}}^{(n)} \delta_{x_n}(x)$

Proof. The proof involves invoking the representer theorem, using bilinearity of scalar products and the fact that if f is a standard Gaussian process then $\forall g \in \mathcal{H} : \langle f, g \rangle \sim \mathcal{N}(0, \|g\|_{\mathcal{H}})$:

$$\mathbb{V}[Z_{f,p}|f_{x_1}, \dots, f_{x_N}] = \quad (24)$$

$$= \mathbb{E}_{f \sim \text{GP}} \left(\int f(x)p(x)dx - \sum_{n=1}^N w_{\text{BQ}}^{(n)} f(x_n) \right)^2 \quad (25)$$

$$= \mathbb{E}_{f \sim \text{GP}} \left(\int \langle f, \phi(x) \rangle p(x)dx - \sum_{n=1}^N w_{\text{BQ}}^{(n)} \langle f, \phi(x_n) \rangle \right)^2 \quad (26)$$

$$= \mathbb{E}_{f \sim \text{GP}} \left\langle f, \int \phi(x)p(x)dx - \sum_{n=1}^N w_{\text{BQ}}^{(n)} \phi(x_n) \right\rangle^2 \quad (27)$$

$$= \|\mu_p - \mu_{q_{\text{BQ}}}\|_{\mathcal{H}}^2 \quad (28)$$

$$= \text{MMD}^2(p, q_{\text{BQ}}) \quad (29)$$

□

We know that the the posterior mean $\mathbb{E}_{\text{GP}}[Z_{f,p}|f_1, \dots, f_N]$ is a Bayes estimator and has therefore the minimal expected squared error amongst all estimators. This allows us to further rewrite ϵ_{BQ}^2 into the following minimax forms:

$$\epsilon_{\text{BQ}}^2 = \sup_{\substack{f \in \mathcal{H} \\ \|f\|_{\mathcal{H}}=1}} \left| \int f(x)p(x)dx - \sum_{n=1}^N w_{\text{BQ}}^{(n)} f(x_n) \right|^2 \quad (30)$$

$$= \inf_{\hat{Z}: \mathcal{X}^N \mapsto \mathbb{R}} \sup_{\substack{f \in \mathcal{H} \\ \|f\|_{\mathcal{H}}=1}} \left| Z - \hat{Z}(f_{x_1}, \dots, f_{x_N}) \right|^2 \quad (31)$$

$$= \inf_{\mathbf{w} \in \mathbb{R}^N} \sup_{\substack{f \in \mathcal{H} \\ \|f\|_{\mathcal{H}}=1}} \left| \int f(x)p(x)dx - \sum_{n=1}^N w_n f(x_n) \right|^2 \quad (32)$$

Looking at ϵ_{BQ}^2 this way, we may discover the deep similarity to the criterion $\epsilon_{\text{herding}}^2$ that kernel herding minimises. Optimal sampling for Bayesian quadrature minimises the same objective as kernel herding, but with the uniform $\frac{1}{N}$ weights replaced by the optimal weights. As a corollary

$$\epsilon_{\text{BQ}}^2(x_1, \dots, x_N) \leq \epsilon_{KH}^2(x_1, \dots, x_N) \quad (33)$$

It is interesting that ϵ_{BQ}^2 has both a Bayesian interpretation as posterior variance under a Gaussian process prior, and a frequentist interpretation as a minimax bound on estimation error with respect to an RKHS.

5 SUBMODULARITY

In this section, we use the concept of approximate submodularity (Krause & Cevher, 2010), in order to study convergence properties of SBQ.

A set function $s : 2^{\mathcal{X}} \mapsto \mathbb{R}$ is *submodular* if, for all $A \subseteq B \subseteq \mathcal{X}$ and $\forall x \in \mathcal{X}$

$$s(A \cup \{x\}) - s(A) \geq s(B \cup \{x\}) - s(B) \quad (34)$$

Intuitively, submodularity is a diminishing returns property: adding an element to a smaller set has larger relative effect than adding it to a larger set. A key result (see e.g. Krause & Cevher, 2010, and references therein) is that greedily maximising a submodular function is guaranteed not to differ from the optimal strategy by more than a constant factor of $(1 - \frac{1}{e})$.

Herding and SBQ are examples of greedy algorithms optimising set functions: they add each pseudosamples in such a way as to minimize the instantaneous reduction in MMD. So it is intuitive to check whether the objective functions these methods minimise are submodular. Unfortunately, neither $\epsilon_{\text{herding}}$, nor ϵ_{BQ} satisfies all conditions necessary for submodularity. However, noting that SBQ is identical to the sparse dictionary selection problem studied in detail by Krause & Cevher (2010), we can conclude that SBQ satisfies a weaker condition called *approximate submodularity*.

A set function $s : 2^{\mathcal{X}} \mapsto \mathbb{R}$ is *approximately submodular* with constant $\epsilon > 0$, if for all $A \subseteq B \subseteq \mathcal{X}$ and $\forall x \in \mathcal{X}$

$$s(A \cup \{x\}) - s(A) \geq s(B \cup \{x\}) - s(B) - \epsilon \quad (35)$$

Proposition 2. $\epsilon_{\text{BQ}}^2(\emptyset) - \epsilon_{\text{BQ}}^2(\cdot)$ is weakly a weakly submodular set function with constant $\epsilon < 4r$, where r is the incoherency

$$r = \max_{x, x' \in \mathcal{P} \subseteq \mathcal{X}} \frac{k(x, x')}{\sqrt{k(x, x)k(x', x')}} \quad (36)$$

Proof. By the definition of MMD we can see that $-\epsilon_{\text{BQ}}^2 = \inf_{w \in \mathbb{R}^N} \left\| \mu_p - \sum_{n=1}^N w_{\text{BQ}}^{(n)} k(\cdot, x_n) \right\|_{\mathcal{H}}^2$ is the negative squared distance between the mean element μ_p and its projection onto the subspace spanned by the elements $k(\cdot, x_n)$. Substituting $k = 1$ into Theorem 1 of Krause & Cevher (2010) concludes the proof. \square

Unfortunately, weak submodularity does not provide the strong near-optimality guarantees as submodularity does. If $s : 2^{\mathcal{X}} \mapsto \mathbb{R}$ is a weakly submodular function with constant ϵ , and $|\mathcal{A}_n| = n$ is the result of greedy optimisation of s , then

$$s(\mathcal{A}_n) \geq \left(1 - \frac{1}{e}\right) \max_{|\mathcal{A}| \leq n} s(\mathcal{A}) - n\epsilon \quad (37)$$

As pointed out by Krause & Cevher (2010), this guarantee is very weak, as in our case the objective function $\epsilon_{\text{BQ}}^2(\emptyset) - \epsilon_{\text{BQ}}^2(\cdot)$ is upper bounded by a constant. However, establishing a connection between SBQ and sparse dictionary selection problem opens up interesting directions for future research, and it may be possible to apply algorithms and theory developed for sparse dictionary selection to kernel-based quasi-Monte Carlo methods.

6 EXPERIMENTS

In this section, we examine empirically the rates of convergence of sequential Bayesian quadrature and herding. We examine both the expected error rates, and the empirical error rates.

In all experiments, the target distribution p is chosen a 2D mixture of 20 Gaussians, whose equiprobability contours are shown in Figure 1. To ensure a comparison fair to herding, the target distribution, and the kernel used by both methods, correspond exactly to the one used in (Chen et al., 2010, Fig. 1). For experimental simplicity, each of the sequential sampling algorithms minimizes the next sample location from a pool of 10000 locations randomly drawn from the base distribution. In practice, one would run a local optimizer from each of these candidate locations, however in our experiments we found that this did not make a significant difference in the sample locations chosen.

6.1 Matching a distribution

We first extend an experiment from (Chen et al., 2010) designed to illustrate the mode-seeking behavior of herding in comparison to random samples. In that experiment, it is shown that a small number of i.i.d. samples drawn from a multimodal distribution will tend to, by chance, assign too many samples to some modes, and too few to some other modes. In contrast, herding places ‘super-samples’ in such a way as to avoid regions already well-represented, and seeks modes that are under-represented.

We demonstrate that although herding improves upon i.i.d. sampling, the uniform weighting of super-samples leads to sub-optimal performance. Figure 1 shows the first 20 samples chosen by kernel herding, in comparison with the first 8 samples chosen by SBQ. By weighting the 8 SBQ samples by the quadrature weights in (20), we can obtain the same expected loss as by using the 20 uniformly-weighted herding samples. Figure 5 shows MMD versus the number of samples added, on the distribution shown in Figure 1. We can see that in all cases, SBQ dominates herding. It appears that SBQ converges at a faster rate than $\mathcal{O}(1/N)$, although the




Figure 5: The maximum mean discrepancy, or expected error of several different quadrature methods. Herding appears to approach a rate close to $\mathcal{O}(1/N)$. SBQ appears to attain a faster, but unknown rate.

form of this rate is unknown.

There are two differences between herding and SBQ: SBQ chooses samples according to a different criterion, and also weights those samples differently. We may ask whether the sample locations or the weights are contributing more to the faster convergence of SBQ. Indeed, in Figure 1 we observe that the samples selected by SBQ are quite similar to the samples selected by kernel herding. To answer this question, we also plot in Figure 5 the performance of a fourth method, which selects samples using herding, but later re-weights the herding samples with BQ weights. Initially, this method attains similar performance to SBQ, but as the number of samples increases, SBQ attains a better rate of convergence. This result indicates that the different sample locations chosen by SBQ, and not only the optimal weights, are responsible for the increased convergence rate of SBQ.

6.2 Estimating Integrals

We then examined the empirical performance of the different estimators at estimating integrals of real functions. To begin with, we looked at performance on 100 randomly drawn functions, of the form:

$$f(x) = \sum_{i=1}^{10} \alpha_i k(x, c_i) \quad (38)$$




Figure 6: Within-model error: The empirical error rate in estimating $Z_{f,p}$, for several different sampling methods, averaged over 250 functions randomly drawn from the RKHS corresponding to the kernel used.

where

$$\|f\|_{\mathcal{H}}^2 = \sum_{i=1}^{10} \sum_{j=1}^{10} \alpha_i \alpha_j k(c_i, c_j) = 1 \quad (39)$$

That is, these functions belonged exactly to the unit ball of the RKHS defined by the kernel $k(x, x')$ used to model them. Figure 6 shows the empirical error versus the number of samples, on the distribution shown in Figure 1. The empirical rates attained by the method appear to be similar to the MMD rates in Figure 5.

By definition, MMD provides a upper bound on the estimation error in the integral of any function in the unit ball of the RKHS (Eqn. (3)), including the Bayesian estimator, SBQ. Figure 7 demonstrates this quickly decreasing bound on the SBQ empirical error.

6.3 Out-of-model performance

A central assumption underlying SBQ is that the integrand function belongs to the RKHS specified by the kernel. To see how performance is effected if this assumption is violated, we performed empirical tests with functions chosen from outside the RKHS. We drew 100 functions of the form:

$$f(x) = \sum_{i=1}^{10} \alpha_i \exp\left(-\frac{1}{2}(x - c_i)^T \Sigma_i^{-1}(x - c_i)\right) \quad (40)$$

where each α_i c_i Σ_i were drawn from broad distributions. This ensured that the drawn functions had features such as narrow bumps and ridges which would




Figure 7: The empirical error rate in estimating $Z_{f,p}$, for the SBQ estimator, on 10 random functions drawn from the RKHS corresponding to the kernel used. Also shown is the upper bound on the error rate implied by the MMD.

not be well modelled by functions belonging to the isotropic kernel defined by k . Figure 8 shows that, on functions drawn from outside the assumed RKHS, relative performance of all methods remains similar.

Code to reproduce all results is available at <http://mlg.eng.cam.ac.uk/duvenaud/>

7 DISCUSSION

7.1 Choice of Kernel

Using herding techniques, we are able to achieve fast convergence on a Hilbert space of *well-behaved* functions, but this fast convergence is at the expense of the estimate not necessarily converging for functions outside this space. If we use a characteristic kernel (Sriperumbudur et al., 2010), such as the exponentiated-quadratic or Laplacian kernels, then convergence in MMD implies weak convergence of q_N to the target distribution. This means that the estimate converges for any bounded measurable function f . The speed of convergence, however, may not be as fast.

Therefore it is crucial that the kernel we choose is representative of the function or functions f we will integrate. For example, in our experiments, the convergence of herding was sensitive to the width of the Gaussian kernel. One of the major weaknesses of kernel methods in general is the difficulty of setting kernel




Figure 8: Out-of-model error: The empirical error rates in estimating $Z_{f,p}$, for several different sampling methods, averaged over 250 functions drawn from outside the RKHS corresponding to the kernels used.

parameters. A key benefit of the Bayesian interpretation of herding and MMD presented in this paper is that it provides a recipe for adapting the Hilbert space to the observations $f(x_n)$. To be precise, we can fit the kernel parameters by maximizing the marginal likelihood of Gaussian process conditioned on the observations. Details can be found in (Rasmussen & Williams, 2006).

7.2 Computational Complexity

While we have shown that Bayesian Quadrature provides the optimal re-weighting of samples, computing the optimal weights comes at an increased computational cost relative to herding. The computational complexity of computing Bayesian quadrature weights for N samples is $\mathcal{O}(N^3)$, due to the necessity of inverting the Gram matrix $K(x, x)$. Using the Woodbury identity, the cost of adding a new sample to an existing set is $\mathcal{O}(N^2)$. For herding, the computational complexity of evaluating a new sample is only $\mathcal{O}(N)$, making the cost of choosing N herding samples $\mathcal{O}(N^2)$. For Monte Carlo sampling, the cost of adding an i.i.d. sample from the target distribution is only $\mathcal{O}(1)$.

The relative computational cost of computing samples and weights using BQ, herding, and sampling must be weighed against the cost of evaluating f at the sample locations. Depending on this trade-off, the three sampling methods form a Pareto frontier over computational speed and estimator accuracy. When computing

method	complexity	rate	guarantee
MCMC	$\mathcal{O}(N)$	variable	ergodic theorem
i.i.d. MC	$\mathcal{O}(N)$	$\frac{1}{\sqrt{N}}$	law of large numbers
herding	$\mathcal{O}(N^2)$	$\frac{1}{\sqrt{N}} \geq \cdot \geq \frac{1}{N}$	(Chen et al., 2010; Bach et al., 2012)
SBQ	$\mathcal{O}(N^3)$	unknown	approximate submodularity

Table 1: A comparison of the rates of convergence and computational complexity of several integration methods.

f is cheap, we may wish to use Monte Carlo methods. In cases where f is computationally costly, we would expect to choose the SBQ method. When f is relatively expensive, but a very large number of samples are required, we may choose to use kernel herding instead. However, because the rate of convergence of SBQ is faster, there may be situations in which the $\mathcal{O}(N^3)$ cost is relatively inexpensive, due to the smaller N required by SBQ to achieve the same accuracy as compared to using other methods.

There also exists the possibility to switch to a less costly sampling algorithm as the number of samples increases. Table 1 summarizes the rates of convergence of all the methods considered here.

8 CONCLUSIONS

In this paper, we have shown three main results: First, we proved that the loss minimized by kernel herding is closely related to the loss minimized by Bayesian quadrature, when selecting sample locations. This implies that sequential Bayesian quadrature can be viewed as an optimally-weighted version of kernel herding.

Second, we showed that the loss minimized by the Bayesian method is approximately submodular with respect to the samples chosen, and established connections to the submodular dictionary selection problem studied in (Krause & Cevher, 2010).

Finally, we empirically demonstrated a superior rate of convergence of SBQ over herding, and demonstrated a bound on the empirical error of the Bayesian quadrature estimate.

8.1 Future Work

In section 5, we showed that SBQ is approximately submodular, which provides only weak sub-optimality guarantees of its performance. It would be of interest to further explore the connection between Bayesian Quadrature and the dictionary selection problem to see if algorithms developed for dictionary selection can provide further practical or theoretical developments. The results in section 6, specifically Figure 5, suggest that the convergence rate of SBQ is faster than

$\mathcal{O}(1/N)$. However, we are not aware of any work showing what the theoretically optimal rate is. It would be of great interest to determine this optimal rate of convergence for particular classes of kernels.

Acknowledgements

The authors would like to thank Carl Rasmussen and Francis Bach for helpful discussions, and Yutian Chen for his help in reproducing experiments.

References

- Bach, F., Lacoste-Julien, S., and Obozinski, G. On the equivalence between herding and conditional gradient algorithms. *Arxiv preprint arXiv:1203.4523*, 2012.
- Chen, Y., Welling, M., and Smola, A. Super-samples from kernel herding. UAI, 2010.
- Cowles, Mary K. and Carlin, Bradley P. Markov chain monte carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904, 1996. ISSN 01621459. doi: 10.2307/2291683. URL <http://dx.doi.org/10.2307/2291683>.
- Krause, A., Guestrin, C., Gupta, A., and Kleinberg, J. Near-optimal sensor placements: Maximizing information while minimizing communication cost. In *Proceedings of the Fifth International Conference on Information Processing in Sensor Networks (IPSN '06)*, pp. 2–10, Nashville, Tennessee, USA, 2006.
- Krause, Andreas and Cevher, Volkan. Submodular dictionary selection for sparse representation. In Fürnkranz, Johannes and Joachims, Thorsten (eds.), *ICML*, pp. 567–574. Omnipress, 2010. ISBN 978-1-60558-907-7.
- Minka, T. P. Deriving quadrature rules from Gaussian processes. Technical report, Statistics Department, Carnegie Mellon University, 2000.
- O’Hagan, A. Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 29:245–260, 1991.
- Rasmussen, C. E. and Ghahramani, Z. Bayesian monte carlo. In Becker, S. and Obermayer, K. (eds.), *Ad-*

vances in Neural Information Processing Systems, volume 15. MIT Press, Cambridge, MA, 2003.

Rasmussen, C.E. and Williams, CKI. Gaussian Processes for Machine Learning. *The MIT Press, Cambridge, MA, USA*, 2006.

Song, Le, Zhang, Xinhua, Smola, Alex, Gretton, Arthur, and Schölkopf, Bernhard. Tailoring density estimation via reproducing kernel moment matching. In *Proceedings of the 25th international conference on Machine learning*, ICML '08, pp. 992–999, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-205-4. doi: 10.1145/1390156.1390281. URL <http://doi.acm.org/10.1145/1390156.1390281>.

Sriperumbudur, Bharath K., Gretton, Arthur, Fukumizu, Kenji, Schölkopf, Bernhard, and Lanckriet, Gert R.G. Hilbert space embeddings and metrics on probability measures. *J. Mach. Learn. Res.*, 99:1517–1561, August 2010. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1859890.1859901>.

Welling, M. Herding dynamical weights to learn. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 1121–1128. ACM, 2009.